arXiv:1410.4296v2 [cs.NI] 20 Oct 2014

Can SDN Mitigate Disasters?

Vincent Gramoli,[*,1,2] Guillaume Jourjon,[1] Olivier Mehani[1]

Abstract: Datacenter networks and services are at risk in the face of disasters. Existing fault-tolerant storage services cannot even achieve a nil recovery point objective (RPO) as client-generated data may get lost before the termination of their migration across geo-replicated datacenters. SDN has proved instrumental in exploiting application-level information to optimise the routing of information. In this paper, we propose *Software Defined Edge (SDE)* or the implementation of SDN at the network edge to achieve nil RPO. We illustrate our proposal with a fault-tolerant key-value store that experimentally recovers from disaster within 30s. Although SDE is inherently fault-tolerant and scalable, its deployment raises new challenges on the partnership between ISPs and CDN providers. We conclude that failure detection information at the SDN-level can effectively benefit applications to recover from disaster.

Keywords: SDN, distributed system

[*]Corresponding author: `vincent.gramoli@sydney.edu.au`
[1] Nicta, Sydney. Eveleigh, NSW, Australia, `first.last@nicta.com.au`
[2] University of Sydney, Sydney, NSW, Australia, `vincent.gramoli@sydney.edu.au`

Copyright © 2014 NICTA
ISSN: 1833-9646-XXXX

1 Introduction

With the advent of cloud services, the computation needed by individuals gets progressively centralised at few datacenters around the world. This centralisation puts the services at risk in the face of disasters, like nuclear power plant explosions or earthquakes, that can affect a large geographical region. This risk motivated, several years ago, some Wall Street financial institutions to build datacenters outside the blast radius of a nuclear attack on New York City, hence drawing a ring of land in New Jersey called the "Donut".[1] Within a range of 30 to 70 km from the city centre, these backup datacenters aim at maintaining synchronous data transfer with the city centre while providing service continuity despite disasters.

This idea as well as all actions taken to ensure service continuity in the event of a disaster are called *disaster recovery* [2]. Recovering some services can be particularly challenging. Take a fault-tolerant key-value store, which serves `get` and `put` requests from clients, as an example. To guarantee that data persist despite the failure of isolated servers, the effect of an update, say a successful `put` request, should at least be replicated at multiple servers of a datacenter. Such a fault-tolerant approach is however not sufficient to recover the data upon a disaster, as all servers of this datacenter could be affected. Prevalent solutions investigated asynchronous migration and mirroring techniques across datacenters of distant geographical locations [9, 15]. The challenge is then to minimise the migration delay as this always translates into some amount of data that can be lost during a disaster, also known as *Recovery Point Objective (RPO)*.

While industrials have made remarkable efforts in reaching low RPO, a disaster often leads to network outage, preventing remote clients from accessing the running backup service during a large delay, referred to as the *Recovery Time Objective (RTO)*. More specifically, if a backup server starts rapidly operating using a different IP address, it may not be instantaneously accessible as refreshing the DNS table at the edge of the network could take hours or even days. The main problem is that the network itself takes a long time for recovering from a disaster, hence delaying the application recovery. Actually, introducing disaster recovery capabilities at the network level should intuitively help achieving both nil RTO and RPO. The challenge lies in distributing the network control that is usually centralised in the network core—typically in the datacenter network—to the network edge.

In this paper, we propose a solution to disaster recovery at the level of the network edge by exploiting *Software Defined Network (SDN)* [8], namely the decoupling of control functions from the data processing and forwarding functions remotely controllable. With the recent advancement in switch fabrics, major content delivery network providers started rerouting the traffic depending on congestion and type of network services. An example is the use of SDN technology by Google for speeding up video streaming content as part of the YouTube service [14]. SDN technology is already available in off-the-shelf network equipment from vendors such as CISCO, Juniper or NEC. More specifically, our solution is to adapt routers at the network edge (1) to prepare for a disaster by proactively replicating the traffic of "critical" data towards multiple datacenters and (2) to cope with a disaster by reactively redirecting the traffic from a failed datacenter to another datacenter located in a non impacted region.

We illustrate our solution by deploying a distributed key-value store application on top of an emulation of our SDN edge or *Software Defined Edge (SDE)* for short. A distributed key-value store is a popular storage abstraction offered by many cloud service providers as part of MongoDB, IBM's Spinnaker, Amazon's SimpleDB, Cassandra or WindowsAzure Table Storage. The popularity of this storage abstraction is in part due to its simplicity in associating a key to a value for simple storage and retrieval of data. Our key-value store is strongly consistent and tolerates isolated failures by exploiting intra-datacenter replication but relies on our SDE solution to cope with disasters. This application is an ideal use-case to illustrate that SDE can reach a nil RPO and a 30s RTO by using TCP retransmissions directly at the router level upon the detection of edge failure and until traffic redirection between datacenter regions as far as Australia to Ireland.

Our SDE solution is not a panacea but raises new kinds of challenges in the deployment of SDN technologies at large scale. An interesting observation is that implementing SDE seems to require a tight collaboration between content delivery network providers and ISPs so that the latter offer an API that the former could exploit to control forwarding rules and mitigate disasters. We are not aware of any such partnership and there is certainly a long way before SDE could be deployed at the scale of Internet. We also believe however that there already exist incentives for such a collaboration. In particular,

[1] http://www.datacenterknowledge.com/archives/2008/03/10/new-york-donut-boosts-nj-data-centers/

we conjecture that SDE would reduce the operational cost of network forensics of these companies as unidirectional traffic such as SYN packets would no longer occurs towards offline datacenter, and thus we may reduce their search range for potential misbehaviours, such as SYN attacks or other DDOS attacks.

The remainder of the paper is organised as follows. In Section 2, we present a brief overview of the related work in disaster recovery and the use of SDN in this particular context. We then introduce our fault-tolerant distributed key-value store in Section 3. In Section 4, we present our SDE architecture, its emulation and the performance results we obtained. Section 5 discusses the opportunities and challenges of our solution. Finally, Section 6 concludes.

2 Related Work

In this section, we present an overview of the storage literature with a particular focus on disaster recovery and how Software Defined Network has been used to cope with scalability and fault tolerance. As far as we know, no SDN solution currently addresses disasters, mainly because existing work focus on privately owned networks or datacenter networks.

2.1 Disaster recovery and storage

The dramatic impact of disasters on the persistence of data has long motivated database research community to investigate how to efficiently cope with disasters [2]. This topic has regained attention in the last few years, with the ermergence of centralised cloud storage services within large datacenters around the globe [15].

In the literature, two metrics usually characterise the persistence of data upon disaster. First, the Recovery Point Objective (RPO) represents the acceptable amount of data that can be lost, while the Recovery Time Objective (RTO) defines the amount of downtime that is permissible before the system recovers. Existing replication techniques usually minimise service latency overhead by executing asynchronously, they periodically copy the freshly stored data or a virtual machine image. Asynchronous replication was shown effective in practice to limit the bandwidth cost [9] even though it cannot achieve a nil RPO. Several companies already offer cloud storage solutions with a reasonably low, yet non-null, RPO. In contrast with asynchronous replication, synchronous replication consists of duplicating the update request so that the two copies remain consistent [15], hence promising to achieve a nil RPO.

Replication is the key to make data resilient to failures. A distributed storage uses replication to guarantee that requests issued by clients get served despite the crash of a server. In the case of a disaster that affects a whole datacenter, it is important to geo-replicate data by copying the data across datacenters located in different regions of the globe. Database solutions traditionally classify servers into a single primary and multiple backups. In such context, they exist two ways of recovering from disaster recovery [2]. First, the client contacts the primary and the primary exchange messages with the backup(s) before the clients gets a response. To recover the data after a disaster at least one backup should be located in a different region from the primary, which makes this solution, called *2-safety*, slow as the client request latency increases with the distance between regions. Second, the client contacts only the primary to get a faster response before any message exchange with the backups, an efficient alternative called *1-safety* that unfortunately cannot guarantee recovery.

Our solution falls between these two approaches. It requires the client to differentiate regular traffic that uses a 1-safety-like approach from the critical traffic that uses a 2-safety-like approach. The critical traffic typically consists of update requests that aim at storing critical data, like financial transactions, that should not be lost despite disasters. By contrast, the regular traffic consists of read-only requests the server does not have to keep track of and normal update requests that are simply guaranteed to survive isolated failures but should complete rapidly. Given that latencies to datacenters vary significantly[2], such a distinction is important to only experience high latency during critical requests.

An alternative to primary/backups is to group servers by majorities, by simply redirecting any request issued by a client to a quorum of the servers. This strategy experiences a longer delay for read-only requests as a client must always wait for the participation of a number of servers that is linear in the

[2]RTT between Oregon and Singapore Amazon EC2 regions is 10× the RTT from Oregon to California Amazon EC2 regions as reported at `http://www.bailis.org/blog/communication-costs-in-real-world-networks`

total number of servers before getting an acknowledgement from the storage service. This is typically the technique used to synchronise remote servers based on the Paxos consensus protocol. Instead of considering majorities, an alternative is to exploit *quorums*, or mutually interesting sets, of servers that indicate the minimum amount of servers where the data should be replicated [10]. While our solution is similar to the combination of 1-safety and 2-safety approaches, it actually balances the load on quorums of servers rather than using the primary/backups approach.

As far as we know, our approach is the first one to reach a $30s$ RTO and a nil RPO, it lets the network recover rapidly from a disaster to make it almost transparent to the storage application.

2.2 Software defined networks

Software defined network (SDN) has been popularised by the decoupling of control functions and the data processing and forwarding functions, which makes data forwarding rules remotely controllable. OpenFlow is one way of defining the control over a collection of input/output ports (the switch) and a flow table[3] that was successfully deployed in the GENI network in the US, the JGN2plus network in Japan and the Fed4FIRE project in Europe.

SDN can be used to adapt networks upon failures. This is especially appealing for datacenters where most of the network-dependent cloud services tend to be gathered. This has led to a large body of work on the topic of datacenter networking [13]. At a lower level, SDN offers inherent fault tolerance capabilities. For example, it allows a controller to use link-layer discovery protocol (LLDP) to detect link failures in a timely fashion, and to adjust in response to such a detection the behavior of the switches it controls. While this early approach may not be adequate for scalability [5], recent versions of OpenFlow even allow the switch to react in face of failures without the continuous help from the controller [11].

The use of SDN-enabled routers is also appealing for offering a cost-effective alternative to feature rich routers by moving the complexity from the switch itself to a separate controller, hence allowing an application-sensitive traffic without embedding the associated complexity at the switch level. Our solution uses a similar concept to distinguish critical traffic from non-critical traffic and duplicate the flow towards remote datacenters.

While SDN can be used as a mean for fault-tolerance, a controller is not immune to failures. It is clear that replication of controllers is not only necessary to scale to a large network but also to allow some controllers to take over the role of failed controllers. The use of multiple controllers was already shown as effective in reducing the fault-tolerance of a particular SDN using NOX controllers and the Mininet virtualised environment [6]. However, disasters typically impact a large region and may potentially affect a large amount of geo-localized network infrastructure equipments, including controllers.

Our solution also guarantees scalability and tolerance to isolated failures by multiplying controllers to forward traffic. The novelty of our SDE approach lies, however, in distributing controllers at the edge of the network to even tolerate disasters affecting potentially an entire privately owned or datacenter network.

3 The Key-Value Store Use-Case

While appealing, the replication needed for fault-tolerance also raises problems related to the consistency of data, by potentially introducing out-of-date replicas or having different versions of the same data hosted at geographically distant datacenters. Hence communication must occur between servers to ensure that the new value of a data updated by a client gets propagated to multiple servers. There are various forms of consistency provided by key-value stores whose strongest is strong consistency or atomicity [4]. Some guarantee eventual consistency tolerating transient inconsistencies while others, like Yahoo's PNUTS, provides another form of consistency, called "timeline", among replicas hosted at remote datacenters.

To make sure of the uniqueness of the value of a data, two concurrent updates of the same data, say `put(k,v)` and `put(k,v')` requests, should be consistently ordered by all servers. This can be achieved using one timestamp per value. This timestamp is computed based on the unique IP address of the issuing client and a global counter that together "tag" the version of a value, indicating for example that `v'` is more up-to-date than `v`. If needed, any server can potentially retrieve and compare the tag

[3] http://archive.openflow.org

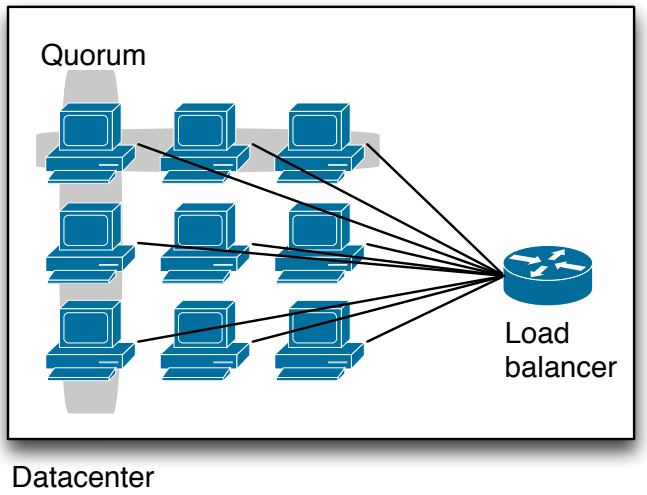

Figure 1: Quorum-based replication within a datacenter

associated to a value to totally order the versions, hence deciding uniformly upon the most recent value of a data.

Another consistency requirement is that an update issued after a previous update on a same key has completed should always have a value tagged with a larger version. This requires that each request, whether it be a read-only, like a `get`, or an update, like a `put`, starts by asking a majority of servers about the current highest version before proceeding with propagating the most up-to-date value with its most up-to-date version. Primary/backups, majorities and quorums are all sets of servers sharing similar intersection properties generalised by the notion of *biquorum system*, corresponding to two classes of server sets such that any set of the first class, say *read-qorums*, always intersect all sets of the second class, say *write-quorums*. One example is the rows and columns of servers as depicted in Figure 1. Using biquorum system, the server receiving the client `get` or `put` request from the client, first obtains the highest value-version pair from a read-quorum (this is the the value with the highest version among the ones received from each server of the quorum) before propagating the highest value-version pair to a write-quorum [1] (making sure that all servers of this quorum stores it locally), such that any read-quorum has at least one server in common with any write-quorum.

To tolerate general failures our approach is twofold. First, we replicate all data within a datacenter where communication is low. This guarantees that the data persists despite isolated failures. Second, we replicate the critical traffic to a dacacenter towards a second datacenter located in a different region. This is the responsibility of the client to distinguish normal from critical data as cross-regions traffic induces significant delays. Provided that both datacenters share a common initial state, say as given by some common virtual machine image, replicating the critical traffic across datacenters guarantees that critical data are not lost upon geo-localized disaster, even if an entire datacenter goes down. We present the architecture to redirect and replicate traffic across distant datacenters in the next section.

4 Redirection of Traffic with SDE

In this section we present a simple yet scalable architecture using SDN technologies at both end of the network to implement a distributed disaster-tolerant key-value store. In particular we propose to deploy SDN technology as close as possible to the client in order to transparently modify the service destination and thus achieving our goal of nil RPO and 30s RTO.

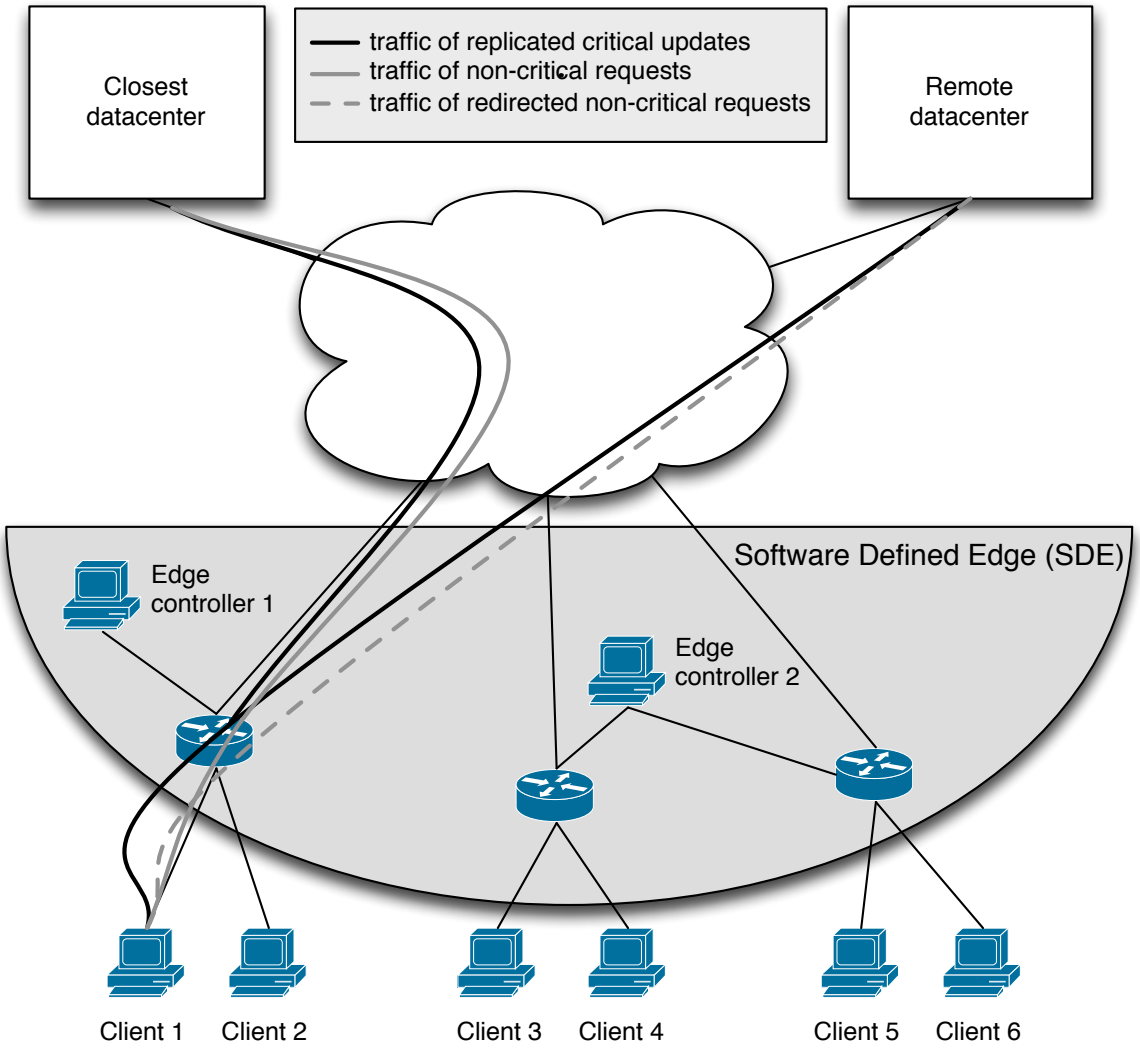

Figure 2: The Software Defined Edge (SDE) replicates critical traffic and redirects (non-critical) traffic upon disaster

4.1 Architecture of the Software Defined Edge

The overall architecture is presented in Figure 2 with a simplified case of a single client connected to the key-value store service deployed over two datacenters. This client can perform two kinds of actions, namely critical update requests, and non-critical requests to the key-value store. In Figure 2, these two actions are symbolised by respectively the black and grey connections. Dashed lines depict the traffic automatically redirected by our architecture upon disaster towards the remote backup datacenter.

In order to deploy our presented architecture, the edge routers on the client side are SDN-capable and connected to an SDN controller. This controller is responsible for the overall disaster resilience and we will detailed its behaviour in the remainder of this section. Furthermore, we also deploy a controller at the entrance of the datacenter in order to balance the load and increase the overall throughput.

For non-critical requests like `get` requests, we propose to use a rapid redirection mechanism. This mechanism is triggered only when the switch controller has detected that the primary datacenter is disconnected from the network by observing that clients cannot connect to their local datacenter. Once the redirection is decided and the controller sends a message to the switches it controls, these switches change the IP headers within the corresponding TCP flow to redirect it to the secondary datacenter

and change back this header on return packets from the datacenter. From the client point of view the connection to the datacenter never changes and appears identical even in the occurrence of a disaster in the region of the primary service provider. We detail this mechanism in the Section 4.2.

In order to decide when to start the redirection, we propose a mechanism inside the edge controller for the detection of potential disaster in the region of the main datacenter. This algorithm takes information from the various edge switches while the decision is centrally taken by the SDN controller responsible of these switches. This centralised algorithm is presented in Algorithm 1 and aims at minimising the RTO depending on the number of services currently deployed.

```
Data: Client flow to service
Result: Possible detection of disaster
detection;
if packet duplicate packet from client then
    possibleDisaster ++ ;
    if possibleDisaster ≥ Threshold then
        Disaster mode = True;
        for every connections affected by disaster do
            send RST packet;
        end
    end
else
    if packet from Service then
        possibleDisaster = 0;
    end
end
```

Algorithm 1: Detection of disaster

4.2 Fast Redirection upon Network Failures

The idea is to place the control plane at the edge to cope with geo-localized disasters. To this end, we use Mininet and the Python-based POX controller to evaluate the performance of our geo-replicated key-value store in the face of disaster.

In this section, we present preliminary results of the proposed architecture depicted in Figure 2 obtained using the mininet [3] software and using POX controller. As a result we have deployed our solution over the topology presented in Figure 3. This topology depicts a local Amazon EC2 datacenter in Sydney and a remote Amazon EC2 datacenter located in Ireland: the link from the switch s_2 to s_3 has a one-way delay of 5ms and a capacity of 100Mbit/s whereas the link from s_2 to s_4 has a one-delay of 170ms and a capacity of 100Mbit/s as the RTT between Sydney and Ireland Amazon EC2 datacenters is known to be about 340ms. Inside each datacenter, we have configured the key-value store over a quorum system of nine servers as depicted in Section 3. Furthermore, both switches s_3 and s_4 act as load balancers inside each datacenter.

Our experiment scenario is as follows. At time $t = 0$, one of the client is starting reading for the local datacenter as fast as possible and thus filling the pipe between s_2 and s_3. At time $t = 200$, we emulate the disaster by putting down the link between s_2 and s_3. Once the disaster occurs, the TCP connection sends retransmission packets and therefore notifies the controller that a possible disaster occurs in the network. As we fixed the detection threshold described in Algorithm 1 to 5 packets, the waiting period varies from 15 to 25 seconds. After this inactive period, the client reconnects to the service with the same IP address but is directed to the second datacenter.

In Figure 4, we present our preliminary results extracted from the captured trace at the client side. The Figure 4a shows the round trip time computed by the TCP connection. In particular, we can see that during the first 200s the flow suffered an RTT of around 100ms, which is induced by the artificial delay we have introduced and the queueing delay when emulating the link capacity of 100Mbit/s. Then, during the final 200s period, the flow experiences the minimal RTT of 340ms necessary to reach the remote datacenter. Figure 4b presents the throughput experienced by the TCP flow. As expected, when the delay is short, the TCP connection is able to fill the pipe when accessing the local datacenter while a

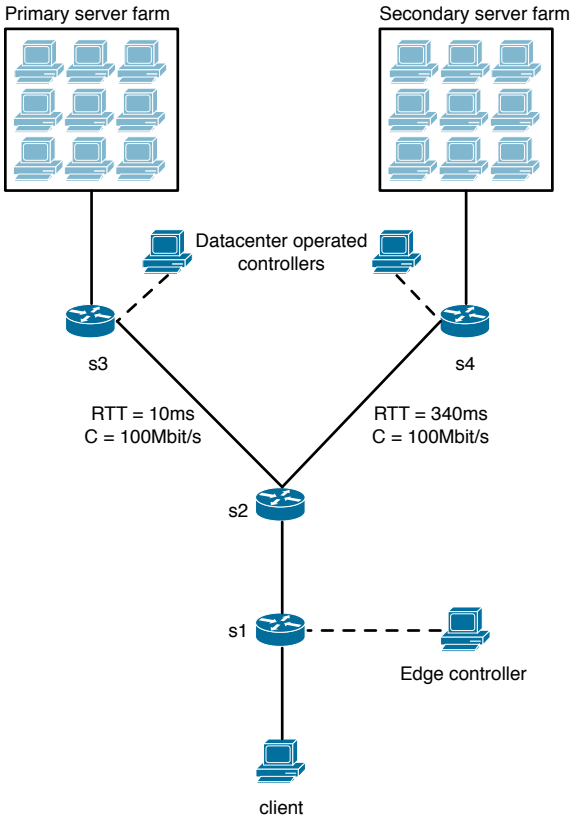

Figure 3: Experimental topology deployed in Mininet to represent a client accessing a local datacenter in Sydney (left) and a remote datacenter in Ireland (right)

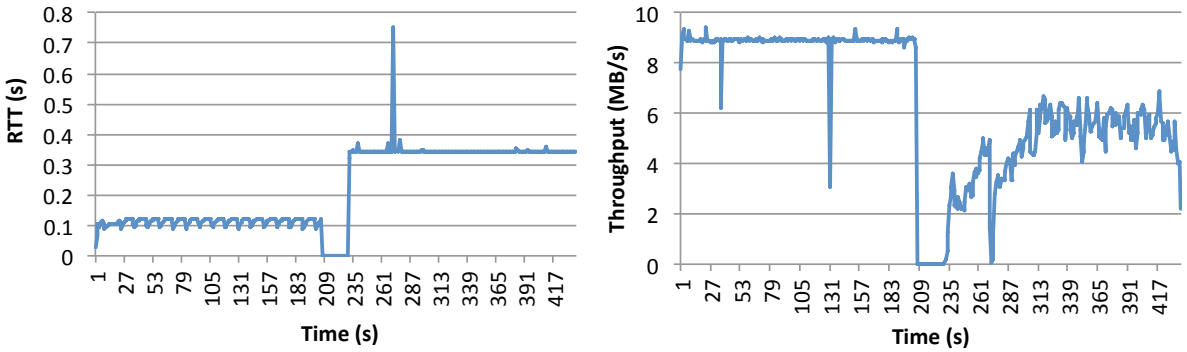

(a) RTT (b) Throughput

Figure 4: Impact of the disaster on service RTT and throughput of the key-value store

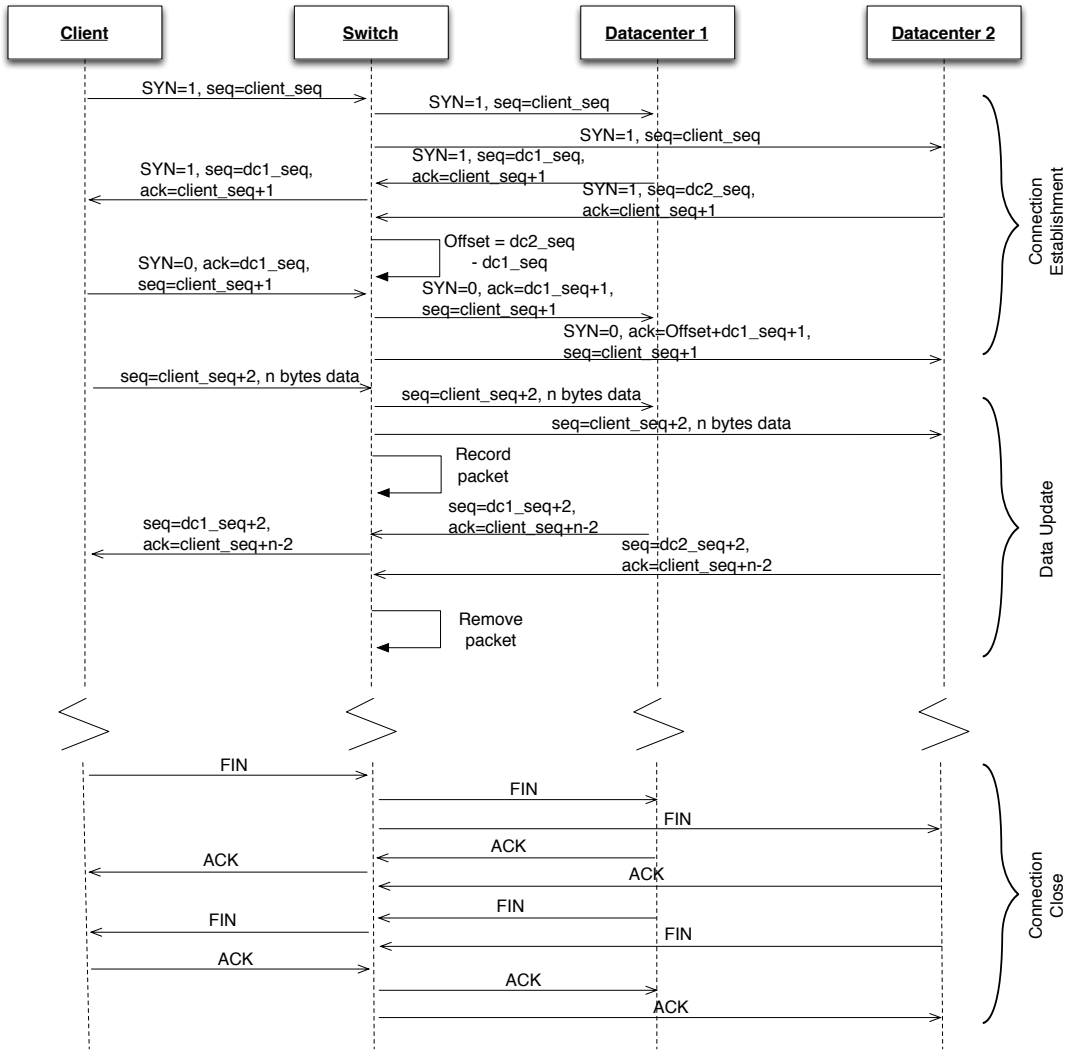

Figure 5: Mechanism for inline TCP duplication

larger RTT makes it impossible for TCP to fully utilise the pipe. Using this SDE technique, we observed empirically an RTO lower than 30s.

4.3 Replication of Traffic with SDE

To ensure persistence of critical data (nil RPO), we extend the redirection mechanism used in the former case by proposing a duplication (or replication) of the TCP flow in the non-disaster period to achieving a 30s RTO. This mechanism is illustrated in Figure 5 where the switch will translate the data from the client to the second datacenter in order to mimic a connection to the client. This replication is made possible by extending the flow table in the switch. In particular, at the establishment of the TCP connection the switch duplicates the SYN packet to both datacenters, while creating an entry in the flow table, then when the first ACK is coming back, the switch elects the datacenter as the master datacenter and keeps a trace of its acknowledged sequence number. Once the second and more distant datacenter sends back the ACK packet, the switch elects this datacenter as the slave and computes the offset in the acknowledged sequence number. This offset is used in the on-going connection to make the slave server believes it is connected to the client. This process is illustrated in the upper part of the sequence diagram of Figure 5.

Once the connection is established between the two servers and the client, the latter can start writing

in the key-value store. In this case, the switch duplicates the TCP stream using the offset value as depicted in Figure 5. As the connection to the second datacenter is longer than to the first datacenter, we need to copy non-acknowledged packets from the slave datacenter in the switch or in the SDN controller. While this imposes a constraint on the storage capability of the switch, we can reasonably expect that the number of flows as well as the size of these flows will not overflow the switch memory given that our solution is deployed at the edge of the network. Using this copy of non-acknowledged packets, we can retransmit them in case of losses in the network.

Finally, we close the client connection to the datacenter in a similar manner as the connection establishment as shown in Figure 5.

5 Challenges and Opportunities

We illustrated how SDE can be used to mitigate disasters. Our proposal faces a number of challenges but also creates new economically viable opportunities. We identified three main challenges as well as three main opportunities and we discuss both aspects in the remainder of this section.

- The first challenge is related to the technical aspect of TCP flow duplication as we mentioned briefly in Section 4.3. Indeed, TCP transmission provides end-to-end reliability that naturally extends to the client to master datacenter connection but the replicated slave connection is beyond the reach of the client. The switch should therefore take care of the retransmission of lost packets based on acknowledgements from the slaves. In the previous section, we introduced an extension to the flow table to perform this action but further investigation is needed to decide when retransmissions are needed. A few options are available, such as timer-based triggers or using a threshold in the number of unacknowledged packets in the buffer. Another issue with TCP duplication concerns the asymmetry of the paths to the datacenters and its impacts on the sending rate.

- Second, the proposed architecture relies on its deployment at the edge of the network. However, it is unlikely that the datacenter operator would be the same as the edge ISP. Hence, the deployment requires collaboration between these two entities. This is a well-known issue for Quality of Service as some argued for the the establishment of an API to delegate some control over the edge switches to third parties [12].

- Finally, we demonstrated our proposal over a simple key-value store. Extension to more complex data storage backends might require additional work.

Notwithstanding these challenges, our proposed architecture presents several opportunities.

- First, we have illustrated how our proposal achieves nil RPO while approaching a nil RTO. This opens a new business opportunity as this feature could be integrated in today's Service Level Agreements as a new premium option.

- Second, while deployment requires collaboration between ISPs and cloud-based service providers, the use of SDN technology greatly reduces its cost. Indeed, SDN-enabled switches are already available off the shelf and usually cost orders of magnitude less than more complex equipments.

- Finally, by proactively redirecting traffic to online datacenter upon disaster, our proposal would reduce the number of legitimate but unsuccessful connection establishments. This in turn would simplify the task of the edge ISP forensic department as a large number of failed connection establishments due to a disaster would have triggered DDoS investigation procedures. Furthermore, we could imagine that the redirection information would be shared with the ISP, which could therefore take the necessary steps to reconfigure their network faster (*e.g.*, change BGP announces).

6 Concluding Remarks

SDN is becoming a mature technology to adapt the control plane within few seconds by dynamically mapping the distribution of switches to controllers [7]. This reactivity promises to facilitate application

fault-tolerance through highly responsive network failure detection. We have confirmed this idea through the use of a simple cloud key-value store application that can recover from continent-wide disaster almost transparently from the end-user point-of-view. Although inter-datacenter SDN is well known to adequately leverage the application information for traffic differentiation, we have shown somehow a reverse ability for distributed applications to leverage SDN edge information about fault detection to effectively bypass disasters.

References

[1] Gregory Chockler, Seth Gilbert, Vincent Gramoli, Peter M. Musial, and Alexander A. Shvartsman. Reconfigurable distributed storage for dynamic networks. *Journal of Parallel and Distributed Computing*, 69(1):100–116, Jan 2009.

[2] Hector Garcia-Molina, Christos A. Polyzois, and Robert B. Hagmann. Two epoch algorithms for disaster recovery. In *Proceedings of the 16th International Conference on Very Large Data Bases*, VLDB '90, pages 222–230, San Francisco, CA, USA, 1990. Morgan Kaufmann Publishers Inc.

[3] Nikhil Handigol, Brandon Heller, Vimal Jeyakumar, Bob Lantz, and Nick McKeown. Reproducible network experiments using container-based emulation. In *CoNEXT*, 2012.

[4] Maurice Herlihy and Jeannette Wing. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3), 1990.

[5] James Kempf, Elisa Bellagamba, András Kern, David Jocha, Attila Takács, and Pontus Sköldström. Scalable fault management for openflow. In *ICC*, pages 6606–6610, 2012.

[6] Hyojoon Kim, J.R. Santos, Y. Turner, M. Schlansker, J. Tourrilhes, and N. Feamster. Coronet: Fault tolerance for software defined networks. In *Proceedings of the 20th IEEE International Conference on Network Protocols*, ICNP, pages 1–2, 2012.

[7] Dan Levin, Andreas Wundsam, Brandon Heller, Nikhil Handigol, and Anja Feldmann. Logically centralized?: State distribution trade-offs in software defined networks. In *HotSDN*, pages 1–6, 2012.

[8] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. Openflow: Enabling innovation in campus networks. *SIGCOMM Comput. Commun. Rev.*, 38(2):69–74, March 2008.

[9] R. Hugo Patterson, Stephen Manley, Mike Federwisch, Dave Hitz, Steve Kleiman, and Shane Owara. SnapMirror: File-system-based asynchronous mirroring for disaster recovery. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*, FAST '02, Berkeley, CA, USA, 2002. USENIX Association.

[10] Shriram Rajagopalan, Brendan Cully, Ryan O'Connor, and Andrew Warfield. SecondSite: Disaster tolerance as a service. In *Proceedings of the 8th ACM SIGPLAN/SIGOPS Conference on Virtual Execution Environments*, VEE '12, pages 97–108, 2012.

[11] Mark Reitblatt, Marco Canini, Arjun Guha, and Nate Foster. FatTire: Declarative fault tolerance for software-defined networks. In *HotSDN*, pages 109–114, 2013.

[12] Vijay Sivaraman, Tim Moors, Hassan Habibi Gharakheili, Dennis Ong, John Matthews, and Craig Russell. Virtualizing the access network via open APIs. In *CoNEXT*, pages 31–42, 2013.

[13] Arsalan Tavakoli, Martin Casado, Teemu Koponen, and Scott Shenker. Applying NOX to the datacenter. In *8th ACM Workshop on Hot Topics in Networks*, HotNets. ACM SIGCOMM, 2009.

[14] Amin Vahdat. Scale and programmability in Google's software defined data center WAN. In *ACM Symposium on Cloud Computing*, SoCC '13, 2013. Personal communication.

[15] Timothy Wood, H. Andrés Lagar-Cavilla, K. K. Ramakrishnan, Prashant Shenoy, and Jacobus Van der Merwe. PipeCloud: Using causality to overcome speed-of-light delays in cloud-based disaster recovery. In *ACM Symposium on Cloud Computing*, SoCC '11, pages 17:1–17:13, 2011.